MS-43459

Information fusion of standoff and other information for biodefense decision support[1]

Jerome J. Braun*[a], Austin Hess[a], Yan Glina[a], Edward C. Wack[a], Karianne Bergen[a], Timothy J. Dasey[a], Robert M. Mays[b], John Strawbridge[b]

[a] Lincoln Laboratory, Massachusetts Institute of Technology
244 Wood Street, Lexington, Massachusetts 02420-9185, USA

[b] JBSDS, Joint Product Manager, Biological Detection Systems,
ATTN: SFAE-CBD-BD-BDS, Bldg. E3548,
Aberdeen Proving Ground, MD 21010-5424, USA

Officially Approved for Public Release by JPEO-CBD

ABSTRACT

This paper discusses selected aspects of an MIT Lincoln Laboratory effort developing information fusion techniques for biodefense decision-support tasks, involving biological standoff (lidar – *light detection and ranging*) sensors, meteorology, as well as point sensors and potentially other battlespace sensing and contextual information. The Spatiotemporal Coherence (STC) fusion approach developed in this effort combines phenomenology aspects with approximate uncertainty measures to quantify corroboration between the information elements. The results indicate that STC can significantly reduce false alarm rates. Meandering Plume and Background Simulation is one of two techniques developed for ground-truth data generation. Beyond the detection realm, developed techniques include information-fusion based plume mapping and propagation prediction.

Keywords: Information fusion, spatiotemporal coherence, biodefense, standoff, detection, mapping, prediction.

1 INTRODUCTION

Decision-makers involved in biological, or chemical, defense are faced with a variety of decisions they have to make accurately and in a timely fashion both before, during and after an attack. Decisions that need to be made prior to any biological attack include resource planning such as sensor placement and utilization, prophylaxis, decontamination, and other options and contingencies. In some situations there might be much time to make them; in others the decisions have to be made quickly, increasing the need for robust decision-support means. The decisions that need to be made once a biological attack has occurred include recognizing attack occurrence, attack characterization, response options, and others. The decision tasks for biodefense decision-support systems include those of detection, plume mapping, propagation prediction (forecasting), identification, course-of-action guidance, adversary intent prediction, and consequence management. Robust decision-support systems can greatly assist making these decisions, especially considering the time pressures faced by the decision-maker.

---

[1] This work was sponsored by the Department of the Army under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

* Email: jbraun@LL.mit.edu

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
2010		**00-00-2010 to 00-00-2010**

4. TITLE AND SUBTITLE **Information fusion of standoff and other information for biodefense decision support**	5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **MIT Lincoln Laboratory,244 Wood St,Lexington,MA,02420-9108**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S) 11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
Tllis paper discusses selected aspects of an MIT Linco~l Laboratory effort developing infonnation fusiou techniques for biodefense decision-support tasks, involving biological standoff (lidar - light detection and ranging) sensors meteorology, as well as point sensors and potentially other battiespace sensing and contextual infomlation. The Spatiotemporal Coherence (STC) fusion approach developed in this effort combines phenomenology aspects with approximate uncertainty measures to qnantify corroboration between the infonnation elements. The results indicate that STC can significantly reduce false alann rates. Meandering Plume and Background Simulation is one of two techniques developed for ground-truth data generation.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			
unclassified	**unclassified**	**unclassified**	**Same as Report (SAR)**	**12**	

Biological sensors are the obvious and necessary "front end" for a biodefense decision-support system. Biological sensing technologies include standoff sensors and point sensors. Biological standoff sensors referred to in this work are lidar based, and in this paper the term standoff and lidar will be used interchangeably. Biological standoff and point sensors differ not only in their general characteristics, but also in terms of their respective strengths. For example, point sensors' particular strengths include sensitivity and specificity. However, as the term implies, point sensors provide indications for their particular location. Standoff sensors on the other hand provide spatial coverage. This can be particularly valuable for early warning or for mapping and tracking the plume. These are some of the tradeoffs that arise when biodefense tasks are viewed from the sensing technology viewpoint.

Information fusion – the algorithmic technology domain that aims at a synergistic exploitation of multiple sensors and other information sources – offers the potential to relax these tradeoffs. It can take advantage of respective strengths of the individual sensors or sensing technologies, and improve the overall system performance beyond that of an individual biological sensor.[3,4,5,6,9]

This paper discusses selected aspects of an MIT Lincoln Laboratory effort developing information fusion techniques for biodefense decision-support tasks, involving biological standoff and other information sources. The approaches developed in this effort take advantage of information sources which are not biological sensors but may offer data potentially relevant to biodefense decision-support tasks. Meteorological data are a prominent example of such information sources. For instance, wind direction and speed influence the plume transport and dispersion, and the techniques developed in this effort take advantage of that relationship. In principle, various other sources such as radar or electro-optical ISR sensors, acoustic, seismic, and other sensor and contextual information sources also can be exploitable by the information-fusion process.

Reduction of false alarms in the bioattack detection task context has been an important objective in this effort, and information-fusion techniques developed to accomplish that objective are discussed. Furthermore, techniques for higher-level decision-support tasks, namely fusion-based plume cloud mapping (tracking) and fusion-based plume propagation prediction (forecasting) techniques developed in this effort are discussed as well. These techniques are applicable to a fixed installation protection and other applications in the military or homeland-protection defense realms.

While development of information fusion algorithms has been a key focus, this effort also required development and advances in several other areas that were needed to support the information-fusion algorithms development main thrust. These supporting areas included sensor modeling, threat scenarios, sensing architectures, techniques for generating data for information-fusion algorithms development, and the development of measures of performance for the performance studies.

This paper discusses some, not all, aspects of this relatively extensive effort. Section 2 discusses the techniques developed for the ground-truth and multisource data generation. Techniques developed for the detection task with a particular goal of reducing false alarms are discussed in section 3. Expanding beyond detection to higher-level decision-support tasks, in section 4 we present the approach developed for fusion-based plume mapping, and in section 5 the approach developed for fusion-based plume propagation prediction is presented.

This paper has a companion paper[2]. The two papers complement one another, in part overlapping and in part covering different elements and aspects of the effort. The present paper assumes that the readers are familiar with the CBRNE defense and sensing domain, including areas such as biodefense-related phenomenology and transport and dispersion models, but are not necessarily focused or involved in algorithmic aspects or information fusion. Consequently, the present paper covers the techniques developed for information-fusion based bioattack detection, plume mapping (tracking), and plume propagation prediction (forecasting), essentially at an overview level. The companion paper[2], on the other hand, assumes readers with information-fusion expertise, including advanced topics such as machine learning, and focuses in a greater depth on the detection-task aspects of this effort. However, since those readers may not necessarily be familiar with the biological and chemical defense domain, the companion paper[2] includes some of the basics related to biological and chemical defense and sensing. There are some overlaps between the two papers. In many cases the overlaps are in terms of topics, but may differ in coverage level of depth. For example, the data-generation techniques are covered in both papers, but the level of coverage is somewhat different. Also, both papers discuss the detection task and the Spatiotemporal Coherence (STC) fusion approach introduced in this effort, but the present paper covers them in broad terms while the companion paper[2] includes more detail. While the present paper discusses in more detail the measures of performance introduced for mapping and propagation prediction, the companion paper[2] addresses in more detail the detection-task measures of performance introduced in this effort. Nevertheless, in the interest of readability and clarity, in some instances the discussion, including the language of some paragraphs (e.g., this paragraph), may be similar between the two papers.

2 GROUND TRUTH AND MULTISOURCE DATA

Availability of release and background data is a prerequisite to information-fusion algorithms development and performance studies. Release data can be obtained from measurements of simulant releases, or by computational simulations. Background data can be obtained from background measurements in given environments, and also from simulations. However, obtaining these data presents a number of challenges. Those challenges and the solutions developed in this effort to address them are briefly outlined in this section.

The data required for information-fusion algorithms' development and studies include data from multiple sensors and potentially other sources, and their temporal progressions. For point sensors, this means indications as a function of time, from potentially multiple non-collocated sensors. For standoff sensors in this work, this involves two-dimensional time-dependent data corresponding to the lidar standoff sensor observations of the area within its *field-of-regard*. Data from meteorology or any other information sources are also generally functions of time and space.

Multisource data are thus time-series of some duration. Each such data sequence is referred to in this work as a *case*, either a background case or a release case. The background cases are those in which no release occurs during the case duration. The release cases are those sequences in which a release does occur at some time-point during the sequence.

The above multisource data represent what would be potentially available as input to the decision-support system. Those data reflect, to a level of accuracy dependent on the sources and on the particular situation, the underlying actual physical reality sensed or evaluated by the information sources that provide those data. That underlying reality is referred to as the *ground truth*. While clearly the ground truth is not normally available to systems other than via the information sources as discussed above, performance studies of systems and algorithms generally need some form of ground truth for use as a reference. Thus, information fusion algorithms' development and studies require both the multisource data and the ground truth data. For the work discussed in this paper, the complete ground truth includes time-varying concentrations of biological substances at all points within the area in which the bioattack detection is to be performed, and the case category (background or release).

A purely measurement-oriented approach for obtaining both ground-truth data and multisensor data corresponding to it is possible. Sensors observing a given area serve as multisource inputs to to the fusion process. Other sensors provide outputs that are by definition considered the ground truth. The data from those sensors are excluded from use in the fusion process, but serve as a reference for performance studies. Clearly, this approach requires measurement experiments as well as the sensors suitable for the ground-truth generation use.

Two approaches were pursued in this work as alternatives to the above. One of these alternatives involved the use of sensor data and sensor models. In that approach, available sensor data are appropriately post-processed and enhanced, yielding outputs that are viewed as the ground truth, but are obviously not used as multisensor data. Rather, the multisensor data are subsequently generated by applying sensor models to that ground truth. The sensor models, fulfilling the role of the sensing process, perturb the ground truth by the desired levels of inexactness. The second approach involved generating the ground truth by computational simulation, and applying sensor models to that ground truth to yield multisource data.

A question might be posed how accurate is the ground truth generated by either of the approaches. This is a challenging question, due to the complex and dynamic nature of the observed aerosol phenomena. However, for information fusion development and performance studies, it is not necessary to resolve this question at an individual plume realization level. As long as the data exhibit the characteristics deemed to reasonably approximate an expected behavior of the phenomena of interest such as an aerosol plume, they can be validly used as ground truth. We shall return to this point when discussing our simulation approach below.

2.1 Release Augmentation approach

The objective of the Release Augmentation approach is to increase the utility of available measured data for information fusion work.[1] The needed data include those representing release-in-background cases under some desired background conditions. However, since the simulant release experiments are normally carried out at test sites, the data measured in those experiments represent specific test-site background conditions. Clearly, those conditions may differ from background conditions in other settings.

As for measuring the *backgrounds alone*, any logistics issues notwithstanding, it is certainly possible to measure backgrounds in variety of settings other than test sites, and data from some background data collections are available. The meteorological conditions, in particular the wind direction and speed, are also normally different between test sites and other settings. All of the above imply that the available measured background data and release-in-background data often differ in terms in background and meteorology conditions. This makes them mutually incompatible for use as input to the information fusion algorithms or performance studies.

The Release Augmentation approach we developed alleviates this issue, increasing the utility of available measured data for use as the information-fusion ground truth. The approach involves computational combining of data from release and background measurements. Briefly, the process starts with two-dimensional measured data, such as those originating from a lidar standoff sensor. The obstructions to standoff sensor, known as *hard targets*, are determined. The release data are processed to extract the release plume from test-site background. The process we refer to as plume tapering computationally extends the plume below the threshold of standoff sensor sensitivity. This allows the generation of ground truth independent of that sensitivity threshold. The resulting plume data are then computationally transformed to reconcile them with the given background data. The wind velocities corresponding to the plume and background data drive this transformation process. This yields the ground truth data.

The Release Augmentation approach allows generation of significantly more ground-truth cases than would be available from sensor measurements alone. Furthermore, it leads to greater flexibility since it enables embedding various plumes in various backgrounds. However, ultimately, the Release Augmentation approach depends on the overall amount of available measured data. The logistics and cost of measurement campaigns places practical constraints on sensor data collections. The simulation approach developed in this effort and described in the following subsection provides a path to generate significant amounts of data unconstrained by experimental measurement-related aspects.

2.2 Meandering Plume and Background Simulation

Modeling of aerosol transport and dispersion has traditionally been approached either with the large-ensemble models or with Computational Fluid Dynamics (CFD) models. Large ensemble models are attractive due to their relative computational efficiency. They provide a probabilistic view of the plume behavior. The output of a large-ensemble model may be thought of as representing in a statistical sense a large number of separate individual plumes rather than any given specific plume (we will also interchangeably refer to single plumes as *single realizations*). In the context of information fusion, the utility of large-ensemble models is limited because the spatio-temporal characteristics of a specific plume are not represented in the large ensemble average view. This becomes quite evident when a large-ensemble plume model output, with its distribution-like shape regularity, is compared to the complex shape evolution and meandering of a typical single-realization plume. CFD models generate single-realization plumes and can be considered a gold standard. Unfortunately, CFD models are computationally extremely costly. This severely limits their practicality as an approach for generating significant amounts of cases needed for the information-fusion work.

Meandering Plume and Background Simulation (MPBS), a computationally efficient single-realization simulation approach we developed, alleviates the limitations and issues discussed above. MPBS allows computationally-efficient simulation of numerous individual release-in-background sequences for a variety of release parameters and background conditions. MPBS combines stochastic and physics-based aspects, reminiscent of our earlier simulation work[5] which served as a starting point for MPBS. However, unlike that approach, MPBS includes a robust representation of turbulence effects, and temporal and spatial correlations. This yields simulation results that include the desired turbulence-driven effects. Major components of MPBS include transport and dispersion, wind-field generation, and turbulence estimation. Figure 1a depicts the overall view of MPBS. An example of MPBS results in Figure 1b includes three time snapshots of a continuous point-release. Figure 1c shows four snapshots of a line release in a scene that includes obstacles. In addition to the meandering, folding and stretching effects, this example result also illustrates the effects of obstacles on MPBS plume propagation.

The question of ground-truth fidelity is sometimes raised in connection with simulation approaches, especially when tradeoffs that simplify an approach are involved. For applications in which the simulation objective is to make predictions, e.g., forecast plume concentration of the progressing release at some specific point and time, this certainly is an important question.

However, as we have argued in the previous publications,[5,6,3] the plume simulation requirements for the information fusion work are more relaxed. Each of the cases simulated by MPBS may or may not be traceable to a specific release and to its corresponding set of release parameters. All that is required is that each case constitutes a *physically plausible* single realization release and its progression, and that the complete set of cases reasonably "covers" in a statistical sense the space of plausible releases. Our results indicate that MPBS meets these objectives.

Figure 1: Meandering Plume and Background Simulation (MPBS) and simulation examples

3 INFORMATION FUSION FOR BIOATTACK DETECTION

Information-fusion approaches developed in this effort for the detection task false-alarm reduction included an approach that was introduced and termed Spatiotemporal Coherence (STC) fusion,[1] adaptive extensions to STC, and a machine-learning approach. The following is a general description of STC. Further details and the discussion of adaptive and machine-learning aspects can be found in the companion paper[2].

STC combines phenomenology and uncertainty aspects to quantify the level of corroboration between the different information sources. This includes in particular the biological standoff, point, and meteorology sensors. However, other sensor or source indications can also be accommodated by STC. Consider a set of alarm events issued by standoff and/or point sensors and the meteorology data such as wind velocity sensor indications. STC determines how these indications "fit together" as part of a potential consistent picture of a progressing bioattack. If they can be satisfactorily "explained" by a putative plume progression, STC increases the confidence level in those alarms accordingly. The less the indications corroborate one another, the lesser overall level of confidence. STC discards alarms that are not sufficiently corroborated by other evidence elements, as false. That leads to a selective suppression of alarms, and may result in false-alarm reduction.

The principle of STC is notionally illustrated in Figure 2a assuming a standoff, point, and meteorology sensing architecture, although it is straightforward to recast this illustration in terms of other sensing sources. The figure shows a standoff sensor indicating an alarm at a certain location, followed by a point sensor alarm at another location. If these two alarm indications corroborate one another, i.e., if it is likely that both alarms correspond to a plume that had progressed consistent with the given wind direction and magnitude, the alarm is deemed high confidence. As the above level of corroboration decreases, so does overall alarm confidence level.

STC involves two types of constructs that govern the coherence determination process. They define the level of uncertainty as function of time and space and are referred to coherence functions. As shown in Figure 2b, the temporal coherence profile is sigmoidal – the more separated the indications are in time, the less they are believed to corroborate one another. The spatial coherence functions are Gaussians with varying parameter values. While the increased spatial distance between alarm indications generally lowers their mutual corroboration level, the degree of corroboration is tied to the temporal aspect, as shown by the arrows in Figure 2b and by the varying amplitude and width of the spatial coherence function instances shown in that figure.

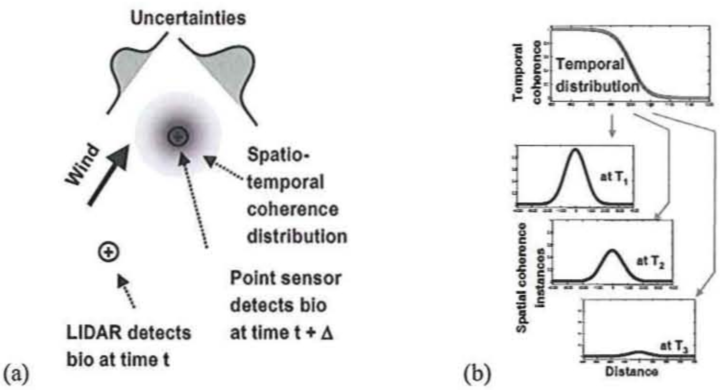

Figure 2: Spatiotemporal Coherence approach: principle and example

A result graph in Figure 3 demonstrates the value of STC for the case of a two biological standoff sensors and meteorology sensing architecture. For this computational experiment, Meandering Plume and Background Simulation (MPBS) data were used to construct a dataset that included multiple release and background cases with varying wind conditions, plume and background conditions. Sensor models that were then activated on that ground truth dataset yielded a multisensor dataset that constituted the input to the STC fusion process.

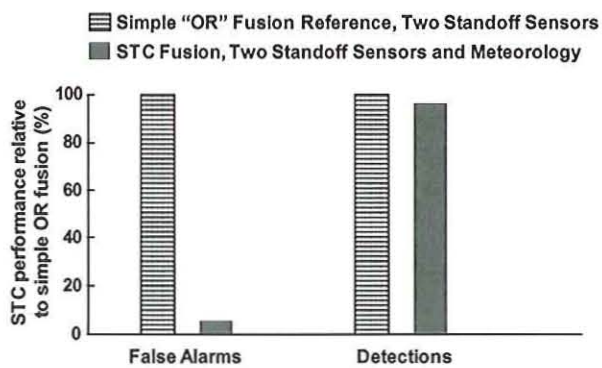

Figure 3: False-alarm reduction using STC fusion approach

Figure 3 compares STC fusion approach involving two standoff sensors and meteorology, to a simple OR fusion approach. Both approaches use standoff sensor alarm indications as inputs, a fusion mode called decision-level fusion. For decision-level fusion, the OR fusion approach constitutes the upper bound on detection, since in OR fusion a detection by any of the two standoff sensors results in a system alarm. However, OR fusion represents the worst-case bound for false alarms, since a false alarm of any of the two standoff sensors results in a system false-alarm. Thus, for decision-level fusion approaches the OR fusion can be useful as a reference to which other approaches such as the STC can be compared. Figure 3 shows that the STC fusion approach offers a competitive detection level, comparable in fact to the OR fusion upper bound, while the false alarms are very substantially reduced. Thus, the STC fusion approach developed in this effort offers a significant potential of false-alarm reduction.

The specifics of the coherence functions, the uncertainty profiles used by STC in the corroboration process, can be determined by the designer. This is one of the manifestations of the intuitive character of STC and constitutes an advantage since human knowledge can be captured in this manner. On the other hand, in some situations this can be a disadvantage. Clearly, a proper selection of the various STC parameters and the associated tuning process requires human involvement and expertise. Furthermore, significant changes in environmental and application specifics, such as changes system deployment settings, may require retuning and consequently an additional human involvement.

Two alternatives to address the above challenge were pursued within this effort. The first involved adaptive provisions for STC, the second involved a supervised learning approach. Both of these are discussed in significantly more detail in the companion paper[2]. In this paper their discussion is limited to the following summary.

The adaptive STC approach involves optimizing STC parameters, such as those related to the coherence function specifics, in an automatic fashion. This requires availability of an optimization dataset of background and release cases. The approach developed in this effort involved a swarm optimization technique. The cost function for the optimization process is the *elliptical distance*, a measure introduced in this effort, which represents both the detection and false-alarm aspects. More specifics on this measure, the optimization process and its results can be found in the companion paper[2]. The machine learning approach that was developed, while retaining some of the STC concepts as features, involves training a supervised machine learning process. The work included feedforward neural networks[10], as well as some Support Vector Machine[8,11] experiments. A more comprehensive discussion of the adaptive STC and the supervised machine-learning approaches can be found in the companion paper[2].

The results obtained within this effort suggest that in some cases additional gains in terms of detections and false-alarm reduction can be achieved with both the adaptive STC and the machine learning alternative. However, the value of these approaches extends beyond that. Both the adaptive STC and the machine learning approaches can take advantage of additional cases that may become available during system operation. In particular, false-alarm cases that may occur can be used to retune the adaptive STC or within a re-training of a machine learning algorithm, in effect adapting the system during its operation. Furthermore, both the optimization procedure in the case of adaptive STC and the training process in case of the machine learning approach can be re-activated as needed, such as in cases of major environmental changes or system application setting changes. There are advantages and tradeoffs between approaches that involve human involvement and tuning such as the STC, and the automatically-tunable approaches such as adaptive STC and the machine learning techniques. The companion paper[2] includes the discussion of these trade-offs and the associated considerations. All of these approaches have advantages and could be used in a hybrid fashion along the lines suggested in our previous publications[3,4]. Whether considered separately or in conjunction, the developed information-fusion techniques for the bioattack detection task, including the STC fusion approach, its adaptive extensions, and the machine-learning approach, offer a significant false-alarm reduction potential.

4 INFORMATION-FUSION BASED PLUME MAPPING

The detection task addresses, as we mentioned in section 1, an important but not the only decision that needs to be made in the context of a bioattack. Once it is determined that an attack has occurred, accurate information regarding the threat plume shape and extent is of significant value to the decision maker. This is referred to as plume mapping. Since the mapping process is performed continuously, the mapping results are a function of time and therefore represent also plume tracking.

From the standpoint of the decision-support and information fusion task hierarchy, in relationship to the detection task the plume mapping process can be viewed as a higher-level task. It should be pointed out that the plume mapping task relies on the detection task and uses detection and discrimination information as one of its inputs.

Since standoff sensors can offer spatial coverage, plume mapping can certainly be performed by a single standoff sensor. In this work the goal was to investigate improving the mapping quality by use of multiple sensors within an information-fusion framework and develop techniques needed for that. Consider a hypothetical situation such as shown in Figure 4. The grey rectangles represent hard targets and the irregular shape at the center of the figure represents the actual plume. The figure is notional, as are the field-of-regard sizes, hard targets and other components shown in it.

In the situation depicted in Figure 4 the entire plume cloud is in the field-of-regard of each of the two standoff sensors, while different hard target obstacles partially obscure it for each of the two. For either of the two sensors in isolation, the obstacles in the sensor's respective field-of-regard constrain the quality of the map obtainable from the sensor data. Using multiple sensors viewing a putative threat plume from different vantage points as shown in Figure 4 may be advantageous. The notional situation shown in Figure 4 is just one of the possible situations in terms of sensors', their fields of regard, hard targets, and plume aspects. For instance, different sensors may have an overlapping view of the plume or they may see separate parts of the plume. Information-fusion techniques that can take advantage of data from multiple sensors that may differ in their field-of-regard or other specifics can lead to improved plume-mapping results.

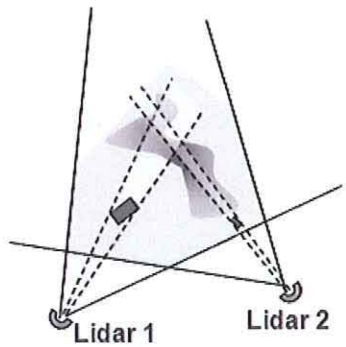

Figure 4: Notional view of one possible configuration of multiple standoff sensors for fusion-based plume mapping

The fusion-based plume mapping approach introduced and developed in this effort extends beyond the use of multiple standoff sensors, as it also uses meteorology data. This allows exploitation of spatio-temporal relationships across different frames corresponding to standoff sensor observations of the plume at different times.

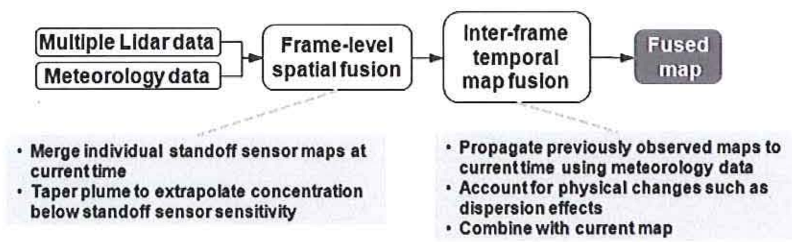

Figure 5: Information fusion-based plume mapping approach

The approach, shown in Figure 5, includes two major stages. The spatial stage fuses the maps corresponding to individual standoff data at a given time. It includes a process introduced in this work and referred to as plume tapering. Conceptually, plume tapering extrapolates concentration data below standoff sensor sensitivity thresholds. The inter-frame temporal map fusion stage fuses map information corresponding to different time-points. More specifically, the previously observed maps are propagated to the current time. That propagation process exploits meteorology, namely the wind direction and wind speed data. The inter-frame temporal map fusion process accounts for physical changes such as the dispersion effects, and fuses the resulting constructs with the current map data. This yields the fusion-based final map for a given time instance. Furthermore, since the process is repeated as the time progresses, yielding a fused map at each time instance, the process in effect also provides plume tracking.

One of the questions that needed to be addressed in context of fusion-based plume mapping was that of measures of performance for mapping. Standard detection-task measures of performance such as ROC curves, or the alternatives we introduced for the detection task as discussed in the companion paper[2], are not suitable for the plume mapping performance studies. Therefore, a new measure of performance was introduced in this work for that purpose. Two aspects must be included in assessing the map quality. Consider a plume and its map produced by some mapping technique. In general, a part of the plume is covered by the produced map while another part of the plume may fall outside of the map. On the other hand, the map may extend beyond the plume covering areas outside it. The former of the two aspects portrays the map relevance in terms of how accurately it maps the plume. The latter portrays the level to which the map over-estimates the actual plume.

The introduced mapping measure of performance includes both of the above aspects, as shown notionally in Figure 6a, and thus can be viewed as a two-component vector measure. One of the components is the percentage of total plume encompassed within the map, the other is the percent of the map that resides outside the plume area. Therefore, it is desired that the map have the value along the first component as high as possible, and along the second component as low as possible. A perfect map will have its first component at 100%, i.e., the entire plume is mapped, while the value for the second will be 0% which corresponds to the absence of any over-estimation. The evaluation of the above measure requires the availability of ground truth, as usual for algorithm performance studies. In addition, in this work a mapping error spatial tolerance margin is used and, since the mapping errors within the tolerance level are neglected, the measure value depends on the specific selection of the tolerance value. Furthermore, a desired contamination threshold must be selected and the actual values of the measure depend on the selected threshold.

Plume mapping performance, whether using a single sensor or multiple sensors, depends on sensor positions, hard targets, and other specifics. Figure 6 illustrates the value of the developed fusion-based mapping approach on a dataset for two standoff sensors. For this dataset, containing multiple release cases, the sensor locations and hard target configuration are identical for all release cases. However, the plumes, wind conditions, and background conditions vary.

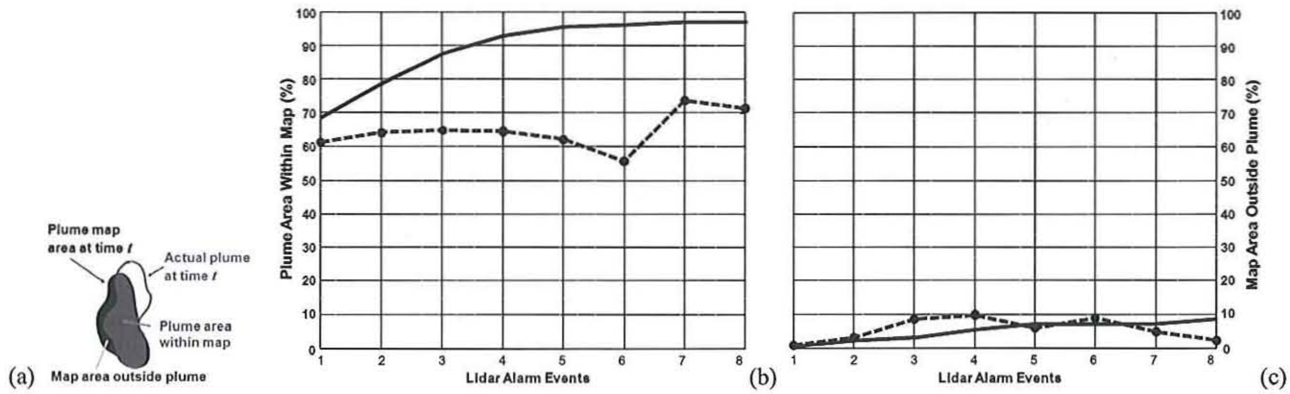

Figure 6: Fusion-based mapping approach results

The fusion-based mapping approach makes use of detection outcomes, and therefore requires and assumes a detection front-end stage. The choice of the detection algorithm used in that front-end stage, and its performance will affect the mapping performance. To gain insight into the efficacy of the fusion-based mapping algorithmic approach it is therefore useful to decouple these two aspects. One way to do that is to investigate the mapping performance under the assumption of an error-free preceding detection stage. The mapping results in Figure 6 were obtained using that assumption. The mapping performance results are shown as a function of the cumulative number of lidar alarm events, which can be interpreted as the amount of standoff information available to the mapping process. The continuous lines correspond to the mapping using the proposed fusion-based approach for the dataset discussed above. The dashed lines correspond to the no-fusion approach which involves standoff data processed in the same way as in the fusion-based approach, but without the multi-standoff and meteorology fusion process. The two components of the vector mapping measure of performance are shown separately as Figure 6b and Figure 6c. However, it should be pointed out that the two are codependent in the sense that, for any given cumulative number of standoff alarm events shown on the horizontal axis, the values of both measure vector components corresponding to that number should be considered simultaneously. (Also, conversely, combining values of two vector components that correspond to different cumulative lidar alarm events number would be meaningless.)

In terms of the graphs in Figure 6, by definition of the mapping vector-measure of performance, it is desirable to attain as high a value as possible in the first graph (Figure 6b), and as low a value as possible in the second (Figure 6c). This corresponds to mapping as much plume area as possible while over-estimating as little as possible. It is evident from Figure 6b that the developed fusion-based mapping approach attains higher values for the first measure-vector component than its no-fusion counterpart. The results shown in Figure 6b are significant in view of the values shown in Figure 6c. That figure shows that the fusion-based mapping approach does not over-estimate appreciably worse than the no-fusion counterpart. Additionally, it is evident from both figures that while with increased amount of lidar evidence the value of measure's first component increases in the fusion-based mapping approach, the second component remains low and essentially unchanged. These results indicate the value of the developed fusion-based plume mapping approach.

5 INFORMATION-FUSION BASED PLUME PROPAGATION PREDICTION

Another higher-level information fusion task is plume propagation prediction. Plume propagation prediction is sometimes referred to as forecasting and, as is well known, transport and dispersion models can be used for that purpose. The first point that should be clarified is how plume propagation prediction task discussed in this section differs from the plume mapping task discussed in section 4. Both tasks aim at providing information related to the plume shape and extent. The major difference is that the mapping task, generating a map at any given time instance t_k, relies on sensor data for that instance t_k, in addition to any potentially-available sensor data for prior instances $t_{i<k}$. That is, the mapping task assumes the availability of a *current* observation. The propagation prediction task, on the other hand, is defined here to start *after* the last standoff-sensor observation and proceeds in absence of any further observations. An example of this is the situation in which the plume that had been observed by a sensors moves beyond their fields of regard. While some specifics of the mapping task discussed in section 4 may be seen as predictive steps, in this work the mapping task is specifically limited to the times of sensor observations mentioned above, and the prediction task assumes prior observations *only*.

The goal of the propagation prediction work in this effort was specifically to investigate the potential of information fusion for the prediction process, not the prediction techniques in general. Specifically, the goal was to develop techniques that may improve the propagation prediction quality by exploiting multi-source information. Two approaches were proposed and developed. Both utilize plume mapping results and meteorology (wind) information as inputs. This is followed by a prediction stage. For both approaches, the starting point is the last available map, i.e., the last observation of the plume by the sensors.

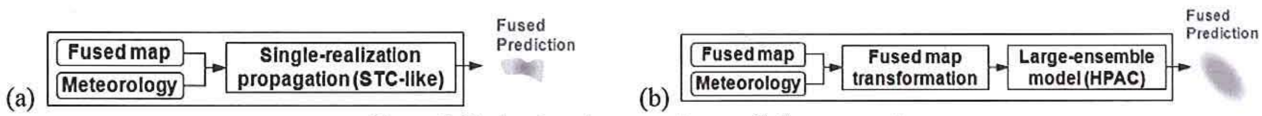

Figure 7: Fusion-based propagation prediction approaches

The first approach, shown in Figure 7a, involves a relatively simple single-realization prediction developed along the principles of the STC approach. Essentially, the map at the time of last plume observation by the sensors is propagated forward using the wind information. Since the last map is a single-realization entity, the approach is single-realization oriented, and the same measure-of-performance used for mapping task can be applied.

The second approach, shown in Figure 7b, involves the use of a large-ensemble model, which in this work was the HPAC model. While the utility of large-ensemble models as a data source for information-fusion efforts is limited, as we discussed in section 2, in context of the fusion-based propagation prediction task their use is appropriate. This is because in the propagation prediction process as defined above there are no additional plume observations the spatiotemporal aspects of which could be exploited. On the contrary, approaching the prediction or forecasting as probabilistic in nature, as is done in large-ensemble models, is well founded.

Embedding a large-ensemble model as part of the fusion-based plume propagation prediction approach is not without challenges. In this effort a novel technique was developed to transform the last-observation maps into forms on which the HPAC model could operate. Also, due to the probabilistic nature of large-ensemble model outputs, performance studies of the propagation prediction approach with an embedded large-ensemble model required introducing yet another measure-of-performance. Comparing a probabilistic distribution-like output of the large-ensemble model such as HPAC to a ground truth of a specific plume realization is inappropriate. This would be tantamount to comparing a single observed value of a variable to its probability distribution. That is generally not meaningful, because a single sample can correspond to any point of the distribution, from its mean value to its tails as in the case of outliers.

The prediction performance studies carried out in this effort involved comparing non-fusion and fusion-based predictions to a reference considered as best prediction. Such best prediction is the one made by the same prediction model when it is activated on the single-realization *ground truth* at the time of the *last* standoff sensor observation (i.e., the last map, either fusion-based or un-fused, from which the prediction process starts). The specifics of the probabilistic prediction measure of prediction performance, and the performance study results which showed the benefit of fusion-based prediction approach are beyond the scope of this paper. Instead, an example in Figure 8 (in this case using the second prediction approach) illustrates the value of the developed fusion-based plume propagation techniques. The fusion-based prediction in this example is more similar to the best-prediction reference than its no-fusion counterpart.

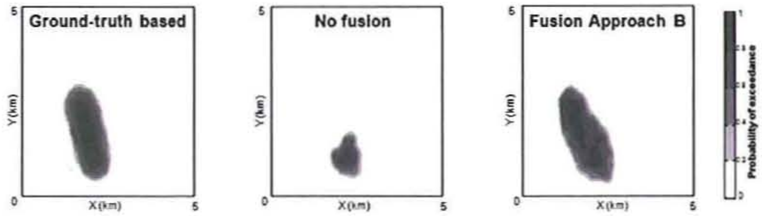

Figure 8: Prediction example

The development of two approaches to fusion-based plume propagation prediction has utility beyond exploring the approaches for their respective benefits or for selecting one of them. These two approaches have different strengths that could be useful dependent on particular specific circumstances faced by the decision-maker. For a given specific prediction case, one of the approaches may produce a more "conservative" prediction, i.e., a larger predicted plume area, while another approach may produce a smaller area. An example of one such case is shown in Figure 9a. In this particular case, the approach that involved an embedded large-ensemble model produced a "more conservative" larger area shown in the figure as the medium-gray area. Selection of that prediction amounts to a lower-risk position assumed by the decision-maker. The STC-like single-realization oriented fusion-based prediction approach offered in this particular case a smaller area, shown as dark-gray area in Figure 9a. Using that prediction rather than the large lower-risk one would amount to the decision-maker taking a more "high-risk" stance. Finally, a relatively simple algorithm was added to generate a medium-risk area that constitutes a simple average-based compromise between the results of the two prediction approaches. The medium-risk area is the light-gray area shown in Figure 9a. If selected, that could correspond to assuming a "medium-risk" tolerance position. Tradeoffs such as these are shown notionally in Figure 9b. The decision-maker can exercise his/her judgment or comfort level to select either of the prediction results as the plume area.

Moreover, multiple approaches such as the two developed in this effort could be incorporated within a decision-support system. Information fusion techniques could rank multiple prediction outcomes, for instance by fusing contextual data that might be available to derive a "current risk level" that would drive the ranking process. The ranked predictions and/or the top-ranking prediction could be offered to the commander or other decision-maker as recommended for the given particular situation and time. This can be viewed as another manifestation of the potential advantages of hybrid paradigms.[3]

a) b)

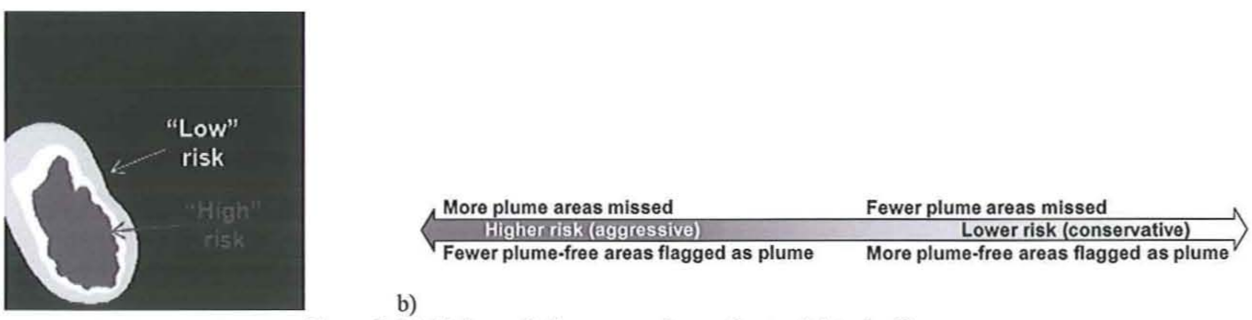

Figure 9: Multiple prediction approaches and potential tradeoffs

6 CONCLUSION

This paper discussed selected aspects of an MIT Lincoln Laboratory effort developing information fusion techniques for biodefense decision-support tasks, involving lidar-based biological standoff sensors, meteorology, as well as point sensors and potentially other battlespace sensing and contextual information. The discussion included techniques developed in the following goal areas: generation of data needed for information fusion algorithms, false-alarm reduction for the bioattack detection task, plume mapping, and plume propagation prediction.

Two approaches developed for generation of ground-truth and multisource data were discussed. The Release Augmentation approach allowed increasing the utility of available experimental data for information fusion development. The Meandering Plume and Background Simulation enabled computationally-efficient generation of individual meandering plumes.

Algorithmic approach developed in this effort and referred to as the Spatiotemporal Coherence (STC) fusion combines phenomenology and uncertainty aspects. The potential of STC to reduce false alarms significantly has been shown. A more extensive discussion of STC, its adaptive extensions, a machine-learning approach, and measures of performance introduced in context with the development of those approaches can be found in the companion paper[2].

Approaches and techniques developed for higher-level decision-support tasks were discussed as well along with the new measures of performance introduced for those tasks. The potential of the developed fusion-based techniques to improve both plume mapping and plume propagation prediction (forecasting) has been shown.

REFERENCES

1. J. J. Braun, Y. Glina, A. Hess, T.J. Dasey, E.C. Wack, R.M. Mays, J. Strawbridge, "Information fusion of biological standoff, point, and other Information sources for bioattack detection," *Proc. Chemical and Biological Defense Physical Science and Technology Conf.*, November, 2008.
2. J. J. Braun, A. Hess, Y. Glina, E. C. Wack, K. Bergen, T.J. Dasey, R.M. Mays, J. Strawbridge, "Approaches to information fusion with spatiotemporal aspects for standoff and other biodefense information sources," *Proc. of SPIE*, vol. 7110, 2010.
3. J. J. Braun and Y. Glina, "Hybrid methods for multisource information fusion and decision support," *Proc. of SPIE*, vol. 6242, 2006.
4. J. J. Braun, Y. Glina, L. Brattain, "Fusion of disparate information sources in a hybrid decision-support architecture," *Proc. of SPIE*, vol. 6571, 2007.
5. J. J. Braun, Y. Glina, J. K. Su, "Urban Biodefense Oriented Aerosol Anomaly Detection using Sensor Data Fusion," *Proc. 1st Joint Conf. Battle Mgt. Nuclear, Chem., Biolog. and Radiological Defense*, November, 2002.
6. J. J. Braun, Y. Glina, J. K. Su, T. J. Dasey, "Computational Intelligence in Biological Sensing," *Proc. of the SPIE*, vol. 5416, 2004.
7. J. J. Braun, Y. Glina, D. W. Stein, P. N. Skomoroch, E. B. Fox, "Information fusion and uncertainty management for biological multisensor systems," *Proc. of SPIE*, vol. 5813, 2005.
8. J. J. Braun, "Sensor Data Fusion with Support Vector Machine Techniques," *Proc. of the SPIE*, vol. 4731, 2002.
9. T.J. Dasey and J.J. Braun, "Information Fusion and Response Guidance", *Lincoln Lab. J.*, vol. 17, no.1, 2007.
10. S. Haykin, *Neural networks: a comprehensive foundation,* Prentice-Hall, Inc., Upper Saddle River, N.J., 1999.
11. V. N. Vapnik, *The Nature of Statistical Learning Theory,* Springer-Verlag New York, Inc., New York, NY, 2000.